arXiv:1511.03003v1 [cs.LO] 10 Nov 2015

Undecidable Cases of Model Checking Probabilistic Temporal-Epistemic Logic*

Ron van der Meyden
Manas K Patra

School of Computer Science and Engineering
UNSW Australia

Abstract

We investigate the decidability of model-checking logics of time, knowledge and probability, with respect to two epistemic semantics: the clock and synchronous perfect recall semantics in partially observed discrete-time Markov chains. Decidability results are known for certain restricted logics with respect to these semantics, subject to a variety of restrictions that are either unexplained or involve a longstanding unsolved mathematical problem. We show that mild generalizations of the known decidable cases suffice to render the model checking problem definitively undecidable. In particular, for a synchronous perfect recall, a generalization from temporal operators with finite reach to operators with infinite reach renders model checking undecidable. The case of the clock semantics is closely related to a monadic second order logic of time and probability that is known to be decidable, except on a set of measure zero. We show that two distinct extensions of this logic make model checking undecidable. One of these involves polynomial combinations of probability terms, the other involves monadic second order quantification into the scope of probability operators. These results explain some of the restrictions in previous work.

1 Introduction

Model checking is a verification methodology used in computer science, in which we ask whether a given model satisfies a given formula of some logic. First proposed in the 1980's [CE81], model checking is now a rich area, with a large body of associated theory and well developed implementations that automate the task of model checking. Significant use of model checking tools is made in industry, in particular, in the verification of computer hardware designs.

Model checking developed originally in a setting where the specifications are expressed in a propositional temporal logic, and the systems to be verified are finite state

*Version of September 15, 2015. This is an extended version, with full proofs, of a paper that appeared in TARK 2015. It corrects an error in the TARK 2015 pre-proceedings version, in the definition of mixed-time polynomial atomic probability formulas.

automata. This setting has the advantage of being decidable, and a great deal of work has gone into the development of algorithms and heuristics for its efficient implementation. More recently, the field has explored the extent to which the expressiveness of both the model representations and of the specification language can be extended while retaining decidability of model checking. Extensions in the systems dimensions considered include real-time systems [ACD90], systems with a mixed continuous and discrete dynamic [MNP08], richer automaton models such as push-down automata, machines with first-in-first out queues etc. In the dimension of the specification language, extensions considered include elements of second order logic and specific constructs to capture the richer properties of the systems models described above (e.g. in the real time case the specification language might contain inequalities over time values.)

Model checking for epistemic logic was first mooted in [HV91], and model checking for the combination of temporal and epistemic logic has been developed both theoretically [MS99, EGM07, HM10] and in practice [GM04, LQR09, KNN$^+$08, Eij04]. A variety of semantics for knowledge are known to be associated with decidable model checking problems in finite state systems, in particular, the observational semantics (in which an agent reasons based on its present observation) the clock semantics (in which an agent reasons based on its present observation and the present clock value), and synchronous and asynchronous versions of perfect recall, all admit decidable model checking in combination with quite rich temporal expressiveness [MS99, EGM07, HM10].

Orthogonally, a line of work on probabilistic model checking has considered model checking of assertions about probability and time [RKNP04]. Although one might at first expect this line of work to be closely related to epistemic model checking, in that probability theory provides a model of uncertainty, in fact this area has been concerned not with how subjective probabilities change over time, but with a probabilistic extension of temporal logic. The focus tends to be on the prior probability of some temporal property, or on the probability that some temporal property holds in runs from a current *known* state.

Rather less attention has been given to model checking the combination of subjective probability and temporal expressiveness. Of the semantics for knowledge mentioned above, the clock and synchronous perfect recall semantics are most suited as a basis for model checking subjective probability. (The others suffer from asynchrony, which makes it more difficult to associate a single natural probability space.) Implementations for these semantics presently exist only for a limited set of formulas, in which the full power of temporal logic is not used. For example, results in [HLM11] for model checking the logic of subjective probability (with clock or synchronous perfect recall semantics) and time restricts the temporal operators to have only finite reach into the future, and does not handle operators such as "at all times in the future".

A fundamental reason underlying this is that the problem of model checking probability with a rich temporal expressiveness seems to be inherently complex. Indeed, it requires a solution to a basic mathematical problem, the *Skolem Problem* for linear recurrences, that has stood unsolved since first posed in the 1930's [Sko34]. Consequently, the strongest results on model checking probability and time that encompass the expressiveness required for model checking knowledge and subjective probability state decidability in a way that requires exclusion of an infinite set of difficult instances for which decidability is unresolved. Specifically, [BRS06] shows that a (weak)

monadic second order logic **PMLO**, containing probability assertions of forms such as $\mathtt{Pr}(\phi(t_1, \ldots, t_n)) > c$, in which the t_i take values in the natural numbers, representing discrete time points, is decidable in finite state Markov chains, provided that the rational number c is not in a set H_ϕ depending on ϕ which can taken to be of arbitrarily small non-zero measure. This work leaves open the decidability of the model checking problem for the language in its full generality, in particular, for the values of c in H_ϕ.

Our contribution in this paper is to consider a number of generalizations of **PMLO**, motivated by model checking a logic of time and subjective probability. In particular, our generalizations arise very naturally when attempting to deal with the way that an agent conditions probability on its observations. We show that these generalizations definitively result in undecidable model checking problems. This clarifies the boundary between the decidable and undecidable cases of model checking logics of probability and time.

We begin in section 2 by recalling the definition of *probabilistic interpreted systems* [Hal03], which provides a very general semantic framework for logics of time, knowledge and probability. We work with an instantiation of this general framework in which systems are generated from finite state partially observed discrete-time Markov chains. We define two logics that take semantics in this framework. The first is an extension of the branching time temporal logic **CTL*** to include operators for knowledge and probability, including operators for the subjective probability of agents. The second is a more expressive monadic second order logic that also adds a capability to quantify over moments of time and *finite sets* of moments of time. In this logic, the agent knowledge and probability operators are indexed by a temporal variable. This logic generalizes the logic of [BRS06]. Our logics allow polynomial comparisons of probability terms, as well as comparisons of agent probability terms referring to multiple time points. We argue from a number of motivating applications that this level of expressiveness is useful in potential applications. We show in Section 3 that the monadic second order logic is as least as expressive as our probabilistic extension of **CTL***. Indeed, some apparently mild extensions of **PMLO** suffice for the encoding: the epistemic and subjective probability operators can be eliminated using a universal modality, polynomial combinations of probability expressions, and a more liberal use of quantification than allowed in **PMLO**.

We then turn in Section 4 to an investigation of the model checking problem. Specifically, we show that model checking even very simple formulas about a single agent's probability is undecidable when the agent has perfect recall. A consequence of this result is that an extension of **PMLO** that adds second order quantification into the scope of probability is undecidable.

This suggests a focus on weaker epistemic semantics instead, in particular, the clock semantics. From the point of view of **PMLO**, to express agent's subjective probabilities with respect to the clock semantics requires polynomial combinations of simple global probability terms of the form " the probability that proposition p holds at time t". We formulate a simple class of formulas involving such polynomial combinations, and show that this also has undecidable model checking.

These results show that even simple model checking questions about subjective probability are undecidable, and moreover help to explain some unexplained restrictions on **PMLO** in [BRS06]: these restrictions are in fact necessary in order to obtain

a decidable logic. We conclude with a discussion of future work in Section 5. Related work most closely related to our results is discussed in the context of presenting and motivating the results.

2 Probabilistic Knowledge

We describe in this section the semantic setting for the model checking problem we consider. We model a set of agents making partial observations of an environment that evolves with time. We first present the semantics of the modal logic we consider, following [Hal03], using the general notion of probabilistic interpreted system. Since these structures are not finite, in order to have a finite input for a model checking problem, we derive a probabilistic interpreted system from a partially observed discrete-time Markov chain. This is done in two ways, depending on the degree of recall of the agents. Taking the Markov chain to be finite, we obtain finitely presented model checking problems whose complexity we then study.

2.1 Probabilistic Interpreted Systems

Probabilistic interpreted systems are defined as follows. Let $Agt = \{1, \ldots, n\}$ be a set of agents operating in an environment e. At each moment of time, each agent is assumed to be in some *local* state, which records all the information that the agent can access at that time. The environment e records "everything else that is relevant". Let S be the set of environment states and let L_i be the set of local states of agent $i \in Agt$. A *global* state of a multi-agent system is an $(n+1)$-tuple $s = (s_e, s_1, \ldots, s_n)$ such that $s_e \in S$ and $s_i \in L_i$ for all $i \in Agt$. We write $\mathcal{G} = S \times L_1 \times \ldots \times L_n$ for the set of global states.

Time is represented discretely using the natural numbers $\mathbb{N}$. A *run* is a function $r : \mathbb{N} \to \mathcal{G}$, specifying a global state at each moment of time. A pair (r, m) consisting of a run r and time $m \in \mathbb{N}$ is called a *point*. If $r(m) = (s_e, s_1, \ldots, s_n)$ then we define $r_e(m) = s_e$ and $r_i(m) = s_i$ for $i \in Agt$. If r is a run and $m \in \mathbb{N}$ a time, we write $r[0..m]$ for $r(0) \ldots r(m)$ and $r_e[0..m]$ for $r_e(0) \ldots r_e(m)$. A *system* is a set $\mathcal{R}$ of runs. We call $\mathcal{R} \times \mathbb{N}$ the *set of points* of the system $\mathcal{R}$.

Agent knowledge is captured using a relation of indistinguishability. Two points (r, m) and (r', m') are said to be *indistinguishable to agent i*, if the agent is in the same local state at these points. Formally, we define $\sim_i$ to be the equivalence relation on $\mathcal{R} \times \mathbb{N}$ given by $(r, m) \sim_i (r', m')$, if $r_i(m) = r'_i(m')$. Relative to a system $\mathcal{R}$, we define the set

$$\mathcal{K}_i(r, m) = \{(r', m') \in \mathcal{R} \times \mathbb{N} \mid (r', m') \sim_i (r, m)\}$$

to be the set of points that are, for agent i, indistinguishable from the point (r, m). Intuitively, $\mathcal{K}_i(r, m)$ is the set of all points that the agent considers possible when it is in the actual situation (r, m). A system is said to be *synchronous* if for all agents i, we have that $(r', m') \in \mathcal{K}_i(r, m)$ implies that $m = m'$. Intuitively, in a synchronous system, agents always know the time. Since it is more difficult to define probabilistic knowledge in systems that are not synchronous, we confine our attention to synchronous systems in what follows.

A *probability space* is a triple $\mathbf{Pr} = (W, \mathcal{F}, \mu)$ such that W is a (nonempty) set, called the *carrier*, $\mathcal{F} \subseteq \mathcal{P}(W)$ is a σ-field of subsets of W, called the *measurable* sets in $\mathbf{Pr}$, containing W and closed under complementation and countable union, and $\mu : \mathcal{F} \to [0, 1]$ is a *probability measure*, such that $\mu(W) = 1$ and $\mu(\bigcup_n V_n) = \sum_n \mu(V_n)$ for every countable sequence $\{V_n\}$ of mutually disjoint measurable sets $V_n \in \mathcal{F}$. As usual, we define the conditional probability $\mu(U|V) = \mu(U \cap V)/\mu(V)$ when $\mu(V) > 0$.

Let *Prop* be a set of *atomic propositions*. A *probabilistic interpreted system* over *Prop* is a tuple $\mathcal{I} = (\mathcal{R}, \mathtt{Pr}_1, \ldots, \mathtt{Pr}_n, \pi)$ such that $\mathcal{R}$ is a system, each $\mathtt{Pr}_i$ is a function mapping each point (r, m) of $\mathcal{R}$ to a probability space $\mathtt{Pr}_i(r, m)$ in which the carrier is a subset of $\mathcal{R} \times \mathbb{N}$, and $\pi : \mathcal{R} \times \mathbb{N} \to \mathcal{P}(Prop)$ is an interpretation of some set *Prop* of atomic propositions. Intuitively, the probability space $\mathtt{Pr}_i(r, m)$ captures the way that the agent i assigns probabilities at the point (r, m), and $\pi(r, m)$ is the set of atomic propositions that are true at the point.

We will work with probabilistic interpreted systems derived from synchronous systems in which agents have a common prior on the set of runs. To define these, we use the following notation. For a system $\mathcal{R}$, a set of runs $\mathcal{S} \subseteq \mathcal{R}$ and a set of points $U \subseteq \mathcal{R} \times \mathbb{N}$, define

$$\mathcal{S}(U) = \{r \in \mathcal{S} \mid \exists m : (r, m) \in U\}$$

to be the set of runs in $\mathcal{S}$ passing through some point in the set U. Conversely, for a set $\mathcal{S}$ of runs and a set U of points, define

$$U(\mathcal{S}) = \{(r, m) \in U \mid r \in \mathcal{S}\}$$

to be the set of points in U that are on a run in $\mathcal{S}$. Note that if there exists a constant $k \in \mathbb{N}$ such that $(r, m) \in U$ implies $m = k$, then the relation $r \leftrightarrow (r, k)$ defines a one-to-one correspondence between $\mathcal{S}(U)$ and $U(\mathcal{S})$. In synchronous systems, in which the sets $\mathcal{K}_i(r, m)$ satisfy this condition, this gives a way to move between sets of points considered possible by an agent and corresponding sets of runs.

Suppose that $\mathcal{R}$ is a synchronous system, let $\mathbf{Pr} = (\mathcal{R}, \mathcal{F}, \mu)$ be a probability space on the system $\mathcal{R}$, and let π be an interpretation on $\mathcal{R}$. Intuitively, the probability space $\mathbf{Pr}$ represents a prior distribution over the runs. We assume that for all points $(r, m) \in \mathcal{R} \times \mathbb{N}$ and agents i, we have that $\mathcal{R}(\mathcal{K}_i(r, m)) \in \mathcal{F}$ is a measurable set and $\mu(\mathcal{R}(\mathcal{K}_i(r, m))) > 0$. (This assumption can be understood as saying that, according to the prior, each possible local state $r_i(m)$ of agent i at time m has non-zero probability of being the local state of agent i at time m.) Under this condition, we define the probabilistic interpreted system $\mathcal{I}(\mathcal{R}, \mathbf{Pr}, \pi) = (\mathcal{R}, \mathtt{Pr}_1, \ldots, \mathtt{Pr}_n, \pi)$ such that $\mathtt{Pr}_i$ associates with each point (r, m) the probability space $\mathtt{Pr}_i(r, m) = (\mathcal{K}_i(r, m), \mathcal{F}_{r,m,i}, \mu_{r,m,i})$ defined by

$$\mathcal{F}_{r,m,i} = \{\mathcal{K}_i(r, m)(\mathcal{S}) \mid \mathcal{S} \in \mathcal{F}\}$$

and such that

$$\mu_{r,m,i}(U) = \mu(\, \mathcal{R}(U) \mid \mathcal{R}(\mathcal{K}_i(r, m))\,)$$

for all $U \in \mathcal{F}_{r,m,i}$. Intuitively, because the set of runs $\mathcal{R}(\mathcal{K}_i(r, m))$ is measurable, we can obtain a probability space with carrier $\mathcal{R}(\mathcal{K}_i(r, m))$ by conditioning in $\mathbf{Pr}$. Because of the synchrony assumption there is, for each point (r, m), a one-to-one correspondence between points in $\mathcal{K}_i(r, m)$ and runs in $\mathcal{R}(\mathcal{K}_i(r, m))$. The construction uses this

correspondence to induce a probability space on $\mathcal{K}_i(r,m)$ from the probability space on $\mathcal{R}(\mathcal{K}_i(r,m))$. We remark that under the additional assumption of perfect recall, it is also possible to understand each space $\mathtt{Pr}_i(r,m+1)$ as obtained by conditioning on the space $\mathtt{Pr}_i(r,m)$. See [Hal03] for a detailed explanation of this point.

2.2 Probabilistic Temporal Epistemic Logic

To specify properties of probabilistic interpreted systems, a variety of logics can be formulated, drawing from the spectrum of temporal logics. Our main interest is in a reasoning about subjective probability and time, so we first consider a natural way to combine existing temporal and probabilistic logics. For purposes of comparison, it is also helpful to consider a rather richer monadic second order logic of probability and time, that is closely related to a logic for which some decidability results are known.

We may combine temporal and probabilistic logics to define a logic **CTL*****KP** that extends the temporal logic **CTL*** by adding operators for knowledge and probability. Its syntax is given by the grammar

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid A\phi \mid X\phi \mid \phi U \phi \mid K_i\phi \mid f(P,\ldots,P) \bowtie c$$

$$P ::= \mathtt{Pr}_i(\phi) \mid \mathtt{Prior}_i(\phi)$$

where $p \in \mathit{Prop}$, c is a rational constant, $\bowtie$ is a relation symbol in the set $\{\leq, <, =, >, \geq\}$, and $f(x_1,\ldots,x_k)$ is multivariate polynomial in k variables $x_1,\ldots x_k$ with rational coefficients. Instances of P are called *basic probability expression*. The instances generated from $f(P,\ldots,P)$ are called *probability expressions*, and are expressions of the form $f(P_1,\ldots,P_k)$, obtained by substituting a basic probability expression P_i for each variable x_i in $f(x_1,\ldots,x_k)$. For example,

$$4\mathtt{Pr}_1(p)^5 \cdot \mathtt{Pr}_2(q)^3 + \frac{7}{15}\mathtt{Pr}_1(p)$$

is an instance of $f(P,\ldots,P)$ obtained from $f(x,y) = 4x^5y^3 + \frac{7}{15}x$ by substituting $\mathtt{Pr}_1(p)$ for x and $\mathtt{Pr}_2(q)$ for y.

Intuitively, formula $K_i\phi$ expresses that agent i knows ϕ. The formula $A\phi$ says that ϕ holds for all possible system evolutions from the current situation. The formula $X\phi$ expresses that ϕ holds at the next moment of time. The formula $\phi_1 U \phi_2$ says that ϕ_2 eventually holds, and ϕ_1 holds until that time. The expression $\mathtt{Pr}_i(\phi)$ represents agent i's current probability of ϕ, $\mathtt{Prior}_i(\phi)$ represents agent i's *prior* probability of ϕ, i.e., the agent's probability of ϕ at time 0. The formula $f(P_1,\ldots,P_k) \bowtie c$ expresses that this polynomial combination of current and prior probabilities stands in the relation $\bowtie$ to c. We use standard abbreviations from temporal logic, in particular, we write $F\phi$ for $\mathit{true}U\phi$.

A restricted fragment of the language that may be of interest is the *branching time fragment* in which the temporal operators are restricted to those of the temporal logic **CTL**. That is, X and U are permitted to occur only in combination with the operator A, in one of the forms $AX\phi$, $EX\phi$, $A\phi_1 U \phi_2$, $E\phi_1 U \phi_2$, where we write $E\phi$ as an abbreviation for $\neg A\neg\phi$. We call this fragment of the language **CTLPK**. The motivation

for considering this fragment is that the complexity of model checking is in polynomial time for the temporal logic **CTL**, whereas it is polynomial-space complete for the richer temporal logic **CTL*** [CES86]. The logic **CTLPK** is therefore, *prima facie*, a candidate for lower complexity once knowledge and probability operators are added to the logic.

The semantics of the language **CTL*****KP** in a probabilistic interpreted system $\mathcal{I} = \mathcal{I}(\mathcal{R}, \mathbf{Pr}, \pi)$ is given by interpreting formulas ϕ at points (r, m) of $\mathcal{I}$, using a satisfaction relation $\mathcal{I}, (r, m) \models \phi$. The definition is mutually recursive with a function $[\cdot]_{\mathcal{I},(r,m)}$ that assigns a value $[P]_{\mathcal{I},(r,m)}$ to each probability expression P at each point (r, m). This requires computing the measure of certain sets. For the moment, we assume that all sets arising in the definition are measurable. We show later that this assumption holds in the cases of interest in this paper.

We first interpret the probability expressions at points (r, m) of the system $\mathcal{I}$, by

$$[\mathtt{Pr}_i\phi]_{\mathcal{I},(r,m)} = \mu_{r,m,i}(\{(r', m') \in \mathcal{K}_i(r, m) \mid \mathcal{I}, (r', m') \models \phi\})$$

$$[\mathtt{Prior}_i\phi]_{\mathcal{I},(r,m)} = \mu_{r,0,i}(\{(r', 0) \in \mathcal{K}_i(r, 0) \mid \mathcal{I}, (r', 0) \models \phi\})$$

$$[f(P_1, \ldots, P_k)]_{\mathcal{I},(r,m)} = f([P_1]_{\mathcal{I},(r,m)}, \ldots, [P_k]_{\mathcal{I},(r,m)})$$

The satisfaction relation is then defined recursively, as follows:

1. $\mathcal{I}, (r, m) \models p$ if $p \in \pi(r, m)$
2. $\mathcal{I}, (r, m) \models \neg\phi$ iff not $\mathcal{I}, (r, m) \models \phi$
3. $\mathcal{I}, (r, m) \models \phi_1 \wedge \phi_2$ iff $\mathcal{I}, (r, m) \models \phi_1$ and $\mathcal{I}, (r, m) \models \phi_2$
4. $\mathcal{I}, (r, m) \models A\phi$ if $\mathcal{I}, (r', m) \models \phi$ for all runs r' with $r'[0 \ldots m] = r[0 \ldots m]$,
5. $\mathcal{I}, (r, m) \models X\phi$ if $\mathcal{I}, (r, m + 1) \models \phi$
6. $\mathcal{I}, (r, m) \models \phi_1 U \phi_2$ holds if there exists $k \geq m$ such that $\mathcal{I}, (r, k) \models \phi_2$, and $\mathcal{I}, (r, l) \models \phi_1$ for all l with $m \leq l < k$.
7. $\mathcal{I}, (r, m) \models K_i\phi$ if $\mathcal{I}, (r', m') \models \phi$ for all $(r', m') \in \mathcal{K}_i(r, m)$.
8. $\mathcal{I}, (r, m) \models f(P_1, ..., P_k) \bowtie c$ if $[f(P_1, ..., P_k)]_{\mathcal{I},(r,m)} \bowtie c$.

2.3 Probabilistic Monadic Second Order Logic

Temporal modal logics refer to time in a somewhat implicit way. An alternative approach is to work in a setting with more explicit references to time, by using variables denoting time points. Kamp's theorem [Kam68] establishes an equivalence in the first order case, but by adding second order variables and quantification, one can obtain richer logics, that frequently remain decidable in the monadic case. In this section, we develop a logic in this style for time and subjective probability.

We define the logic **WMLOKP** as follows. We use two types of variables: time variables t and set variables X. Time variables take values in $\mathbb{N}$ and set variables take *finite* subsets of $\mathbb{N}$ as values. *Probability terms* P have the form $\mathtt{Pr}(\phi)$ or the form

$\mathtt{Pr}_{i,t}(\phi)$ where $i \in Agt$ is an agent, t is a time variable, ϕ is a formula. Formulas ϕ are defined by the following grammar:

$$\begin{aligned}\phi ::= \quad & p(t) \mid X(t) \mid t_1 < t_2 \mid f(P, \ldots, P) \bowtie c \mid \neg\phi \mid \phi \wedge \phi \mid \\ & K_{i,t}(\phi) \mid \forall t(\phi) \mid \forall X(\phi)\end{aligned}$$

where t, t_1, t_2 are time variables, p is an atomic proposition, X is a set variable, i is an agent, $c \in \mathbb{Q}$ is a rational constant, f is a rational polynomial (see the discussion above for **CTL*****KP**), and $\bowtie$ is a relation symbol from the set $\{=, <, \leq, >, \geq\}$.

Intuitively, in this logic formulas are interpreted relative to a run. Instead of indexing by a single moment of time, as in the logic above, we relativize the satisfaction relation to an assignment of values to the temporal and set variables. Atomic formula $p(t)$ says that proposition p holds at time t. Similarly, a (finite) set X of times can be interpreted as a proposition, and we can understand $X(t)$ as stating that the value of t is in X. (We remark that there is a fundamental difference between the types of propositions denoted by atomic propositions p and set variables X: whereas the atomic propositions may depend on structural aspects of the run, such as the global state at time t, the set variables may refer only to the time.) The atomic formula $t_1 < t_2$ has the obvious interpretation that time t_1 is less than time t_2. The constructs $\forall t(\phi)$ and $\forall X(\phi)$ correspond to universal quantification over times and *finite* sets of times respectively. They say that ϕ holds on the current run for all values of the variable. (Taking finite sets amounts to the *weak* interpretation of second order quantification. One could also consider a strong semantics allowing infinite sets of times. We have opted here for the weak interpretation to more easily relate our results to the existing literature.)

The probability term $\mathtt{Pr}(\phi)$ refers to the probability of ϕ in the probability space on runs. The meaning of probability term $\mathtt{Pr}_{i,t}(\phi)$ is agent i's probability at time t that the run satisfies ϕ. Similarly, $K_{i,t}\phi$ says that agent i knows at time t that the run satisfies ϕ. Note that, whereas in **CTL*****KP**, the formula $K_i\phi$ always expresses that agent i knows that ϕ holds at the "current time", in **WMLOKP**, formulas such as

$$\exists u(u < t \wedge K_{i,t}(p(u)))$$

talk about the agent's knowledge, at some time t, about what was true at some earlier time u. A similar point applies to probability expressions.

Accordingly, for the semantics of **WMLOKP**, we use a variant of interpreted systems in the form $\mathcal{I} = (\mathcal{R}, \Pr, \pi)$, where $\mathcal{R}$ is a system, i.e., a set of runs, and π is an interpretation, as above, but where $\Pr = (\mathcal{R}, \mathcal{F}, \mu)$ is a probability space with carrier equal to the set of runs $\mathcal{R}$, rather than a mapping associating a probability space over a set of points with each agent at each point.

When dealing with formulas with free time and set variables, we need the extra notion of an assignment for the time and set variables. This is a function τ such that for each free time variable t we have $\tau(t) \in \mathbb{N}$, and for each free set variable X we have that $\tau(X)$ is a finite subset of $\mathbb{N}$. Given such an assignment, we give the semantics of probability terms and formulas by a mutual recursion. We give the semantics of formulas ϕ by means of a relation $\mathcal{I}, \tau, r \models \phi$ defined as follows:

1. $\mathcal{I}, \tau, r \models p(t)$ if $p \in \pi(r, \tau(t))$, when p is an atomic proposition,

2. $\mathcal{I}, \tau, r \models X(t)$ iff $\tau(t) \in \tau(X)$, if X is a set variable,

3. $\mathcal{I}, \tau, r \models t_1 < t_2$ iff $\tau(t_1) < \tau(t_2)$,

4. $\mathcal{I}, \tau, r \models \neg\phi$ iff not $\mathcal{I}, \tau, r \models \phi$,

5. $\mathcal{I}, \tau, r \models \phi_1 \wedge \phi_2$ iff $\mathcal{I}, \tau, r \models \phi_1$ and $\mathcal{I}, \tau, r \models \phi_2$,

6. $\mathcal{I}, \tau, r \models K_{i,t}(\phi)$ if $\mathcal{I}, \tau, r' \models \phi$ for all $(r', m') \in \mathcal{K}_i(r, \tau(t))$,

7. $\mathcal{I}, \tau, r \models f(P_1, ..., P_k) \bowtie c$ if $[f(P_1, ..., P_k)]_{\mathcal{I},\tau,r} \bowtie c$,

8. $\mathcal{I}, \tau, r \models \forall t(\phi)$ if $\mathcal{I}, \tau[t \mapsto n], r \models \phi$ for all $n \in \mathbb{N}$,

9. $\mathcal{I}, \tau, r \models \forall X(\phi)$ if $\mathcal{I}, \tau[X \mapsto U], r \models \phi$ for all finite $U \subseteq \mathbb{N}$.

In item (7), the definition is mutually recursive with the semantics of probability terms, which are interpreted as real numbers, relative to a temporal assignment. We define

$$[\mathtt{Pr}(\phi)]_{\mathcal{I},\tau,r} = \mu(\{r' \in \mathcal{R} \mid \mathcal{I}, \tau, r' \models \phi\})$$

and

$$[\mathtt{Pr}_{i,t}(\phi)]_{\mathcal{I},\tau,r} = \frac{\mu(\{r' \in \mathcal{R} \mid (r, \tau(t)) \sim_i (r', \tau(t)),\ \mathcal{I}, \tau, r' \models \phi\})}{\mu(\{r' \in \mathcal{R} \mid (r, \tau(t)) \sim_i (r', \tau(t))\})}$$

$$[f(P_1, \ldots, P_k)]_{\mathcal{I},\tau,r} = f([P_1]_{\mathcal{I},\tau,r}, \ldots, [P_k]_{\mathcal{I},\tau,r})$$

As above, we assume measurability of the sets required, and also that the agent probability expressions do not involve a division by zero. We later justify that this holds in the particular setting of interest in this paper.

A particular class of formulas of **WMLOKP** will be of interest below. Define a *mixed-time polynomial atomic probability formula* to be a formula of the form[1]

$$\exists t_1 \ldots t_n (f(\mathtt{Pr}(\phi_1), \ldots, \mathtt{Pr}(\phi_m)) = 0)$$

where $f(x_1, \ldots, x_m)$ is a rational polynomial and each ϕ_i is an atomic formula of the form $p(t_j)$ for some proposition p and $j \in \{1 \ldots n\}$. We motivate the usefulness of such temporal mixing of probability expressions in Section 2.5.

The logic **WMLOKP** generalizes several logics from the literature. If we restrict the language by excluding the probability comparison atoms $f(P_1, \ldots, P_k) \bowtie c$ and knowledge formulas $K_{i,t}(\phi)$, we have the *Weak Monadic Logic of Order*, which is equivalent to WS1S [Buc60]. We obtain the *Probabilistic Monadic Logic of Order* considered in [BRS06], which we denote here by **PMLO**, if we

- exclude the knowledge operators $K_{i,t}$,
- exclude agent's probability terms $\mathtt{Pr}_{i,t}(\phi)$, and
- limit the global probability comparisons to be of the form $\mathtt{Pr}(\phi(t_1, \ldots, t_k)) \bowtie c$, containing just a single probability term $\mathtt{Pr}(\phi(t_1, \ldots, t_k))$, with the further constraint that the only free variables of ϕ should be temporal variables $t_1, \ldots t_k$.

[1]The TARK 2015 pre-proceedings version of this paper incorrectly had a universal quantifier in this definition. The existential form is needed for the correctness of Theorem 12.

In particular, second-order quantification into probability expressions, e.g., $\forall X[\mathtt{Pr}(X(t)) > c]$ is not permitted in **PMLO**, but second order quantification that does not cross a probability operator, such as $\mathtt{Pr}(\forall X[X(t)]) > c$, is allowed. We note that **PMLO** *does* allow first order quantifications into the scope of probability, such as $\forall t[\mathtt{Pr}(p(t)) > c]$.

In the sequel, we refer to quantification into the scope of a knowledge formula or probability expression as *quantifying-in*.

2.4 Partially Observed Markov Chains

Although they provide a coherent semantic framework, probabilistic interpreted systems are infinite structures, and therefore not suitable as input for a model checking algorithm. We therefore work with a type of finite model called an *interpreted partially observed discrete-time Markov chain*, or PO-DTMC for short. A finite PO-DTMC for n agents is a tuple $M = (S, PI, PT, O_1, ..., O_n, \pi)$, where S is a finite set of states, $PI : S \to [0..1]$ is a function such that $\sum_{s \in S} PI(s) = 1$, component $PT : S \times S \to [0, 1]$ is a function such that $\sum_{s' \in S} PT(s, s') = 1$ for all $s \in S$, and for each agent $i \in Agt$, we have a function $O_i : S \to O$ for some set O. Finally, $\pi : S \to \mathcal{P}(Prop)$ is an interpretation of the atomic propositions $Prop$ at the states.

Intuitively, $PI(s)$ is the probability that an execution of the system starts at state s, and $PT(s, t)$ is the probability that the state of the system at the next moment of time will be t, given that it is currently s. The value $O_i(s)$ is the observation that agent i makes when the system is in state s. (Below, in the context of interpreted systems, we treat the set of states S as the states of the environment rather than as the set of global states. Agents' local states will be derived from the observations.)

Note that the first three components (S, PI, PT) of a PO-DTMC form a standard discrete-time Markov chain. This gives rise to a probability space on runs in the usual way. A *path* in M is a finite or infinite sequence $\rho = s_0 s_1 \ldots$ such that $PI(s_0) \neq 0$ and $PT(s_k, s_{k+1}) > 0$ for all k with $0 \leq k < |\rho| - 1$. We write $\mathtt{P}_\infty(M)$ for the set of all infinite paths of M, and, for $m \in \mathbb{N}$, write $\mathtt{P}_m(M)$ for the set of all finite paths $s_0 s_1 \ldots s_m$ with exactly m transitions. Any finite path $\rho = s_0 s_1 \ldots s_m$ defines a set

$$\mathtt{P}_\infty(M) \uparrow \rho = \{\omega \in \mathtt{P}_\infty(M) \mid \omega[0 \ldots m] = \rho\} \tag{2}$$

That is, $\mathtt{P}_\infty(M) \uparrow \rho$ consists of all infinite paths which have ρ as a prefix.

We now define a probability space $\mathbf{Pr}(M) = (\mathtt{P}_\infty(M), \mathcal{F}, \mu)$ over the set $\mathtt{P}_\infty(M)$ of all infinite paths of M. The σ-algebra $\mathcal{F}$ is defined to be the smallest σ-algebra over $\mathtt{P}_\infty(M)$ that contains as basic sets all the sets $\mathtt{P}_\infty(M) \uparrow \rho$ for $\rho = s_0 s_1 \ldots s_m$ a finite path of M. For these basic sets, the function μ is defined by

$$\mu(\mathtt{P}_\infty(M) \uparrow \rho) = PI(s_0) \cdot PT(s_0, s_1) \cdot \ldots \cdot PT(s_{m-1}, s_m) \,.$$

The fact that μ can be extended to a measure on $\mathcal{F}$ is a non-trivial result of Kolmogorov for more general stochastic processes [KSK76].

We may construct several different probabilistic interpreted systems from each PO-DTMC, depending on what agents remember of their observations. We consider two, one that assumes that agents have perfect recall of all their observations, denoted `spr`, and the other, denoted `clk`, which assumes that agents are aware of the current time

and their current observation. Recall that runs in an interpreted system map time to global states, consisting of a state of the environment and a local state for each agent. We interpret the states of the PO-DTMC M as states of the environment. To obtain a run, we also need to specify a local state for each agent at each moment of time. We use the the observations to construct the local states.

In the case of the *synchronous perfect recall semantics*, given a path $\rho \in \mathtt{P}_\infty(M)$, we obtain a run $\rho^{\mathtt{spr}}$ by defining the components at each time m as follows. The environment state at time m is $\rho_e^{\mathtt{spr}}(m) = \rho(m)$, and the local state of agent i at time m is $\rho_i^{\mathtt{spr}}(m) = O_i(\rho(0)) \dots O_i(\rho(m))$. Intuitively, this local state assignment represents that the agent remembers all its past observations. We write $\mathcal{R}^{\mathtt{spr}}(M)$ for the set of runs of the form $\rho^{\mathtt{spr}}$ for $\rho \in \mathtt{P}_\infty(M)$. Note that this system is synchronous: if $r = \rho^{\mathtt{spr}}$ and $r' = \omega^{\mathtt{spr}}$ then for each agent i and time $m \in \mathbb{N}$, if $r_i(m) = r'_i(m')$, then $O_i(\rho(0)) \dots O_i(\rho(m)) = O_i(\omega(0)) \dots O_i(\omega(m'))$, which implies $m = m'$.

For the *clock semantics*, we construct a run a $\rho^{\mathtt{clk}}$ in which again the environment state at time m is $\rho_e^{\mathtt{clk}}(m) = \rho(m)$, and for agent i we define the local state at time m by $\rho^{\mathtt{clk}}(m) = (m, O_i(\rho(m)))$. Intuitively, this says that the agent is aware of the clock value and its current observation. We write $\mathcal{R}^{\mathtt{clk}}(M)$ for the set of runs of the form $\rho^{\mathtt{clk}}$ for $\rho \in \mathtt{P}_\infty(M)$ an infinite path of M. This system is also synchronous: if $r = \rho^{\mathtt{clk}}$ and $r' = \omega^{\mathtt{clk}}$ then for each agent i and time $m \in \mathbb{N}$, if $r_i(m) = r'_i(m')$, then $(m, O_i(\rho(m))) = (m', O_i(\omega(m')))$, hence $m = m'$. In both cases of $x \in \{\mathtt{spr}, \mathtt{clk}\}$, if T is a subset of $\mathtt{P}_\infty(M)$, we write T^x for $\{\rho^x \mid \rho \in T\}$.

In both cases of $x \in \{\mathtt{spr}, \mathtt{clk}\}$, we have a one-to-one correspondence between the infinite paths $\mathtt{P}_\infty(M)$ and the runs $\mathcal{R}^x(M)$. We therefore can induce probability spaces $\mathbf{Pr}^x(M)$ on $\mathcal{R}^x(M)$ from the probability space $\mathbf{Pr}(M)$ on $\mathtt{P}_\infty(M)$. As described above, the probability space $\mathbf{Pr}^x(M)$ on runs moreover induces a probability space $\mathtt{Pr}_i^x(r, m)$ on the set of points considered possible by each agent i at each point (r, m). The PO-DTMC M gives us an interpretation π on its states, and we may derive from this an interpretation π^x on the points (r, m) of $\mathcal{R}^{\mathtt{spr}}(M)$ and $\mathcal{R}^{\mathtt{clk}}(M)$ by defining $\pi^x(r, m) = \pi(r_e(m))$. Using the general construction defined above, we then obtain the probabilistic interpreted systems $\mathcal{I}^x(M) = \mathcal{I}(\mathcal{R}^x(M), \mathbf{Pr}^x(M), \pi^x)$ for $x \in \{\mathtt{spr}, \mathtt{clk}\}$.

It is necessary to establish the measurability of certain sets for the semantic definitions of the logics above to be complete. We now establish this for the systems $\mathcal{I}^{\mathtt{spr}}(M)$ and $\mathcal{I}^{\mathtt{clk}}(M)$.

First, the general construction of $\mathcal{I}(\mathcal{R}, \mathbf{Pr}, \pi)$, where $\mathbf{Pr} = (\mathcal{R}, \mathcal{F}, \mu)$, assumed that for all $m \in \mathbb{N}$ and all runs r, we have that $\mathcal{R}(\mathcal{K}_i(r, m)) \in \mathcal{F}$ is measurable and $\mu(\mathcal{R}(\mathcal{K}_i(r, m))) > 0$. The following lemma assures us that this is the case for both the perfect recall and the clock constructions. We write $\mathcal{K}_i^x(r, m)$ for $\mathcal{K}_i(r, m)$, relative to the set of runs $\mathcal{R} = \mathcal{R}^x(M)$.

Lemma 1 *Let M be a finite PO-DTMC. Then for each $x \in \{\mathtt{spr}, \mathtt{clk}\}$, and run $r \in \mathcal{R}^x(M)$, the set $\mathcal{R}^x(M)(\mathcal{K}_i^x(r, m))$ is measurable in $\mathbf{Pr}^x(M)$ and has measure > 0.*

Proof: For $x \in \{\mathtt{spr}, \mathtt{clk}\}$, define the relations $\sim^x$ on finite paths of M as follows. For finite paths $\rho = s_0 \dots s_m$ and $\omega = t_0 \dots t_k$, define $\rho \sim_i^{\mathtt{clk}} \omega$ if $m = k$ and $O_i(s_m) = O_i(t_m)$, and $\rho \sim_i^{\mathtt{spr}} \omega$ if $m = k$ and $O_i(s_j) = O_i(t_j)$ for all $j = 0 \dots m$. Note that in both cases, $\rho \sim_i^x \omega$ implies that $|\rho| = |\omega|$. Since the number of paths of length m is finite, it

follows that the number of ω such that $\rho \sim_i^x \omega$ is finite. If $r \in \mathcal{R}^x(M)$ then there exists $\rho \in \mathtt{P}_\infty(M)$ such that $r = \rho^x$. Let $U = \{\omega \in \mathtt{P}_m(M) \mid \omega \sim_i^x \rho[0 \ldots m]\}$. It is straightforward from the definitions that

$$\mathcal{R}^x(M)(\mathcal{K}_i^x(\rho^x, m)) = \bigcup_{\omega \in U} (\mathtt{P}_\infty(M) \uparrow \omega)^x$$

The sets $\mathcal{R}^x(M)(\mathcal{K}_i^x(\rho^x, m))$, for $x \in \{\mathtt{spr}, \mathtt{clk}\}$, are therefore each a finite union of measurable sets, hence measurable. The fact that their measure is > 0 follows from the fact that both contain $(\mathtt{P}_\infty(M) \uparrow \rho)^x$, which has non-zero measure since ρ is path of M. □

Next, in giving the semantics of **CTL*****KP**, we require that sets of the form $\{(r', m') \in \mathcal{K}_i^x(r, m) \mid \mathcal{I}^x(M), (r', m') \models \phi\}$ are measurable in $\mathbf{Pr}_i^x(r, m)$. Similarly, in the semantics of **WMLOKP**, we require that sets of the form $T_\tau^x(\phi) = \{r \in \mathcal{R}^x(M) \mid \mathcal{I}^x(M), \tau, r \models \phi\}$ are measurable, where τ is a temporal assignment and ϕ is a formula of **WMLOKP**. This is established in the following result.

Lemma 2 *Let M be a finite PO-DTMC and $x \in \{\mathtt{spr}, \mathtt{clk}\}$. For every set $S \subseteq \mathcal{R}(M)$ of runs of M such that the semantic definitions above of* **CTL*****KP** *and* **WMLOKP** *in $\mathcal{I}^x(M)$ refer to $\mu(S)$, the set S is measurable in* $\mathbf{Pr}(M)$.

Proof: We show this just for the logic **WMLOKP**; for **CTL*****KP** the result is a corollary derivable using using the fact that **CTL*****KP** can be translated to **WMLOKP** (Proposition 3 below). We use structural induction on ϕ to show that $T_\tau^x(\phi)$ is measurable for all temporal assignments τ.

If $\phi = q(t)$ where q is an atomic proposition and t is a temporal variable, then note that for all paths $\rho \in \mathtt{P}_{\tau(t)}(M)$ and runs $r \in (\mathtt{P}_\infty(M) \uparrow \rho)^x$, we have $\mathcal{I}^x(M), \tau, r \models q(t)$ iff $q \in \pi(r(\tau(t)))$ iff $q \in \pi(\rho(\tau(t)))$. Thus,

$$T_\tau^x(\phi) = \bigcup_{\rho \in \mathtt{P}_{\tau(t)}(M),\ q \in \pi(\rho(\tau(t)))} (\mathtt{P}_\infty(M) \uparrow \rho)^x$$

which is a finite union of basic measurable sets, hence measurable.

If $\phi = X(t)$ where X is a set variable and t is a temporal variable, then $T_\tau^x(\phi)$ is $\mathcal{R}(M)$ in case $\tau(t) \in \tau(X)$ and is $\emptyset$ otherwise. In either case, it is measurable. Similarly, if $\phi = t_1 < t_2$ then $T_\tau^x(\phi)$ is either $\mathcal{R}(M)$ or $\emptyset$, and therefore measurable.

If $\phi = \phi_1 \wedge \phi_2$ then $T_\tau^x(\phi) = T_\tau^x(\phi_1) \cap T_\tau^x(\phi_2)$ is measurable, because, by the inductive hypothesis, $T_\tau^x(\phi_1)$ and $T_\tau^x(\phi_2)$ are. Similarly, $T_\tau^x(\neg\phi) = \mathcal{R}^x(M) - T_\tau^x(\phi)$ is measurable because $T_\tau^x(\phi)$ is measurable, by induction.

For the quantifiers, we have

$$T_\tau^x(\forall t(\phi)) = \bigcap_{m \in \mathbb{N}} T^x_{\tau[t \mapsto m]}(\phi)$$

and

$$T_\tau^x(\forall X(\phi)) = \bigcap_{U \subset \mathbb{N},\ U\ finite} T^x_{\tau[X \mapsto U]}(\phi)\ .$$

In either case, we have a countable intersection of sets that are measurable, by induction, hence these sets are also measurable.

For the knowledge formula $\phi = K_{i,t}\psi$, note that for $\rho \in \mathtt{P}_\infty(M)$, if $\mathcal{I}^x(M), \tau, \rho^x \models K_{i,t}\psi$, then

$$(\mathtt{P}_\infty(M) \uparrow \rho[0 \ldots \tau(t)])^x \subseteq T^x_\tau(K_{i,t}\psi) \,,$$

else

$$(\mathtt{P}_\infty(M) \uparrow \rho[0 \ldots \tau(t)])^x \cap T^x_\tau(K_{i,t}\psi) = \emptyset \,.$$

This is because, under both the clock and perfect recall semantics, satisfaction of $K_{i,t}\psi$ depends only on the observations made to time $\tau(t)$. Taking

$$U = \{\rho[0 \ldots \tau(t)] \mid \rho \in \mathtt{P}_\infty(M) \text{ and } \mathcal{I}^x(M), \tau, \rho^x \models K_{i,t}\psi\} \,,$$

we have that

$$T^x_\tau(K_{i,t}\psi) = \bigcup_{\rho \in P(i,\tau,\psi)} (\mathtt{P}_\infty(M) \uparrow \rho)^x$$

is a finite union of basic measurable sets, hence measurable.

Finally, consider the probability formulas $\phi = f(P_1, \ldots, P_k) \bowtie c$, where $P_j = \mathtt{Pr}_{i_j,t_j}(\psi_j)$ or $P_j = \mathtt{Pr}(\psi_j)$. Let m be the maximum of the values $\tau(t_j)$ for t_j a temporal variable in the former type of probability term. Note that the value $[\mathtt{Pr}(\psi_j)]_{\mathcal{I},\tau,r}$ is independent of the run r, and if r and r' are runs with $r[0, \ldots, m] = r'[0, \ldots, m]$ then $[\mathtt{Pr}_{i_j,t_j}(\psi_j)]_{\mathcal{I}^x(M),\tau,r} = [\mathtt{Pr}_{i_j,t_j}(\psi_j)]_{\mathcal{I}^x(M),\tau,r'}$. It follows also that when $r[0, \ldots, m] = r'[0, \ldots, m]$, we have $\mathcal{I}^x(M), \tau, r \models f(P_1, \ldots, P_k) \bowtie c$ iff $\mathcal{I}^x(M), \tau, r' \models f(P_1, \ldots, P_k) \bowtie c$. Let

$$U = \{\rho[0, \ldots, m] \mid \rho \in \mathtt{P}_\infty(M),\ \mathcal{I}^x(M), \tau, \rho^x \models f(P_1, \ldots, P_k) \bowtie c\}$$

be the set of all prefixes of length m of paths satisfying $f(P_1, \ldots, P_k) \bowtie c$. The set U is finite, and we have, by the above, that

$$T^x_\tau(f(P_1, \ldots, P_k) \bowtie c) = \bigcup_{\rho \in U} (\mathtt{P}_\infty(M) \uparrow \rho)^x$$

is a finite union of basic measurable sets, hence measurable. □

2.5 Discussion

We have defined our logics to be quite expressive in the type of atomic probability assertions we have allowed, which involve polynomials of probability expressions. In **WMLOKP**, these expressions may explicitly refer to different time points. Some existing logics of probability in the literature use a more restricted expressiveness, e.g., [FH94] consider a logic that has only linear combinations of probability expressions, and many logics [BRS06, RKNP04] allow only inequalities involving a single probability term. Here give some motivation to show that the richness we have allowed is natural and useful for applications.

Polynomials: There are several motivations for allowing polynomial combinations of probability expressions. One, as noted in [FHM90], is that polynomials arise naturally from conditional probability. If we would like to include linear combinations of conditional probability expressions in the language, we find that this motivates a generalization to polynomial combinations of probability expressions. Consider the formula $\mathtt{Pr}(\phi_1|\psi_1) + \mathtt{Pr}(\phi_2|\psi_2) \leq c$. Expanding out the definition of conditional probability, we have

$$\frac{\mathtt{Pr}(\phi_1 \wedge \psi_1)}{\mathtt{Pr}(\psi_1)} + \frac{\mathtt{Pr}(\phi_2 \wedge \psi_2)}{\mathtt{Pr}(\psi_2)} \leq c \,.$$

We see here that there is a risk of division by zero that needs to be managed in order for the semantics of this formula to be fully defined. One way to do so is to multiply out the denominators, resulting in the form

$$\mathtt{Pr}(\phi_1 \wedge \psi_1) \cdot \mathtt{Pr}(\psi_2) + \mathtt{Pr}(\phi_2 \wedge \psi_2) \cdot \mathtt{Pr}(\psi_1) \leq c \cdot \mathtt{Pr}(\psi_1) \cdot \mathtt{Pr}(\psi_2)$$

which is meaningful in all cases. (Should this not have the desired semantics in case one of the $\mathtt{Pr}(\psi_i)$ is zero, an additional formula can be added that handles this special case as desired.) However, although we started with a linear probability expression, we now have multiplicative terms. This suggests that the appropriate way to add the expressiveness of conditional probability to the language is to admit atomic formulas that compare polynomial combinations of probability expressions.

More generally, although it is less of relevance for purposes of model checking, and more of use for axiomatization of the logic, allowing polynomials also naturally enables familiar reasoning patterns to be captured inside the logic. In particular, validities such as $\mathtt{Pr}(\phi_1 \vee \phi_2) = \mathtt{Pr}(\phi_1) + \mathtt{Pr}(\phi_2)$ when ϕ_1 and ϕ_2 are mutually exclusive and $\mathtt{Pr}(\phi_1 \wedge \phi_2) = \mathtt{Pr}(\phi_1) \cdot \mathtt{Pr}(\phi_2)$ when ϕ_1 and ϕ_2 are independent show that both addition and multiplication of probability terms arises naturally.

Mixed-time: A second way in which our logics are rich is in allowing probability atoms that refer to different moments of time. In **CTL*KP** this already the case because combinations such as $\mathtt{Prior}_A(\phi) = \mathtt{Pr}_A(\phi)$ are allowed, which refer to both the current time and to time 0. The logic **WMLOKP** takes such temporal mixing further by allowing reference to time points explicitly named using time variables.

Such temporal mixing is natural, since there are potential applications that require this expressiveness. For example, in computer security, one often wants to say that the adversary A does not learn anything about a secret from watching an exchange between two parties. However, it is often the case that the adversary knows some prior distribution over the secrets. (For example, the secret may be a password, and choice of passwords by users are very non-uniform, with some passwords like '123456' having a very high probability.) This means that the simple assertion that the adversary does not know the secret, or that the adversary has a uniform distribution over the secret, does not capture the appropriate notion of security. Instead, as recognised already by Shannon in his work on secrecy [Sha49], we need to assert that the adversary's distribution over the secret has not changed as a result of its observations. This requires talking about the adversary's probability at two time points. For example, [HLM11] capture an anonymity property by means of formulas using terms $\mathtt{Prior}_A(\phi) = \mathtt{Pr}_A(\phi)$.

Mixed-time polynomials: Additionally, the logic of probability applied to formulas referring to different times leads naturally to polynomial combinations of probabil-

ity terms, each referring to a different moment of time. For example, although **PMLO** allows only formulas of the form $\mathtt{Pr}(\phi(t_1,\ldots,t_n)) \bowtie c$, where the t_i are time variables, the decision algorithm of [BRS06] uses the fact that, when $t_1 < t_2 < \ldots < t_n$, the formula $\phi(t_1,\ldots,t_n)$ is equivalent to a formula of the form $\phi_1(t_1) \wedge \phi_2(t_2 - t_1) \wedge \ldots \wedge \phi_n(t_n - t_{n-1}) \wedge \phi_{n+1}(t_n)$, where the $\phi_i(t)$ are independent past-time formulas for $i = 1 \ldots n$ and $\phi_{n+1}(t)$ is a future time formula. (This statement is closely related to Kamp's theorem [Kam68].) This enables $\mathtt{Pr}(\phi(t_1,\ldots,t_n))$ to be expressed as a sum of products of terms of the form $\mathtt{Pr}(\phi_i(u))$ where $\phi_i(u)$ has just a single free time variable u. Thus, although mixed-time probability formulas are not directly expressible in the logic of [BRS06], specific ones are implicitly expressible, and the extension is a mild one. It is worth remarking, however, that the coefficients of the polynomial expansion of $\mathtt{Pr}(\phi(t_1,\ldots,t_n))$ are all positive, so we do not quite have arbitrary polynomials here. We return to this point below.

3 Relating the logics

The logic **WMLOKP** is very expressive, so it is not surprising that it can capture all of **CTL*****KP**. The following result makes this precise.

For the results below, it is convenient to add to the system a special agent $\bot$ that is blind, and an agent $\top$ that has complete information about the state. In the context of PO-DTMC's these agents are obtained by taking the observation functions to satisfy $O_\bot(s) = O_\bot(t)$ and $O_\top(s) = s$ for all states s, t. We write $\Box\phi$ for $K_{\bot,t}\phi$ where t is any time variable. This gives a *universal modality*: $\Box\phi$ says that ϕ holds on all runs. We write $[t \mapsto n]$ for the temporal assignment defined only on temporal variable t, and mapping this to n.

Proposition 3 *Let M be a PO-DTMC with agent $\top$ and let $x \in \{\mathtt{spr}, \mathtt{clk}\}$. For every formula ϕ of* **CTL*****KP**, *there exists a formula $\phi^*(t)$ of* **WMLOKP** *with t the only free variable, such that $\mathcal{I}^x(M), (r,n) \models \phi$ iff $\mathcal{I}^x(M), [t \mapsto n], r \models \phi^*(t)$ for all runs r.*

Proof: The translation is defined by the following recursion:

$$
\begin{aligned}
&p^*(t) = p(t)\\
&(\neg\phi)^*(t) = \neg\phi^*(t)\\
&(\phi_1 \wedge \phi_2)^*(t) = \phi_1^*(t) \wedge \phi_2^*(t)\\
&(X\phi)^*(t) = \exists u(u = t+1 \wedge \phi^*(u))\\
&(K_i\phi)^*(t) = K_{i,t}(\phi^*(t)),\\
&(\phi_1 U \phi_2)^*(t) = \exists u \geq t(\phi_2^*(u) \wedge \forall v(t \leq v < u \Rightarrow \phi_1^*(v)))\\
&(\mathtt{Pr}_i(\phi))^*(t) = \mathtt{Pr}_{i,t}(\phi^*(t))\\
&(\mathtt{Prior}_i(\phi))^*(t) = \mathtt{Pr}_{i,0}(\phi^*(0))\\
&(f(P_1,\ldots,P_k) \bowtie c)^*(t) = f(P_1^*(t),\ldots,P_k^*(t)) \bowtie c
\end{aligned}
$$

Note that $u = v$ is definable as $\neg(u < v \vee v < u)$, that $u = t + 1$ is definable as $u > t \wedge \forall v > t\ (u \leq v)$, and that $u = 0$ is definable as $\neg\exists t(u = t + 1)$. We can

use $(A\phi)^*(t) = K_{\top,t}(\phi^*(t))$ to translate $A\phi$ in the perfect recall case. In case of the clock semantics, this translation loses the information about the initial state, which is required for correctness of the translation of $\mathtt{Prior}_i(\phi)$. In this case, we introduce, without loss of generality, new propositions p_s for each state s, such that $p_s \in \pi_e(t)$ iff $s = t$, and take

$$(A\phi)^*(t) = \bigwedge_{s \in S} (p_s(0) \Rightarrow K_{\top,t}(p_s(0) \Rightarrow \phi^*(t))) \; .$$

□

With respect to the specific systems we derive from PO-DTMC's with respect to the clock and perfect recall semantics, we are able to make some further statements that simplify the logic **WMLOKP** by eliminating some of the operators. These results are useful for the undecidability results that follow.

For the following results, we note that, without loss of generality, we may assume that a finite PO-DTMC comes equipped with atomic propositions that encode the observations made by the agents. Specifically, when agent i has possible observations $O_i(S) = \{o_{i,1}, \ldots, o_{i,k_i}\}$, we assume that there are atomic propositions $obs_{i,j}$ for $i \in Agt$ and $j = 1 \ldots k_i$ such that for all states s, we have $obs_{i,j} \in \pi(s)$ iff $O_i(s) = o_{i,j}$. Thus, $obs_{i,j}(t)$ holds in a run just when agent i makes observation $o_{i,j}$ at time t.

Proposition 4 *With respect to $\mathcal{I}^{\mathtt{clk}}(M)$ for a finite PO-DTMC M, the operators $K_{i,t}$ and $\mathtt{Pr}_{i,t}$ can be eliminated using the universal operator $\Box$ and polynomial comparisons of universal probability terms $\mathtt{Pr}(\psi)$, respectively. For simple probability formulas $\mathtt{Pr}_{i,t}(\phi) \bowtie c$, only linear probability comparisons are required.*

Proof: The formula

$$\bigwedge_{j=i \ldots k_i} (obs_{i,j}(t) \Rightarrow \Box(obs_{i,j}(t) \Rightarrow \phi)$$

is easily seen to be equivalent to $K_{i,t}(\phi)$ in $\mathcal{I}^{\mathtt{clk}}(M)$. Similarly, $\mathtt{Pr}_{i,t}(\phi) \bowtie c$ can be expressed as

$$\bigwedge_{j=i \ldots k_i} (obs_{i,j}(t) \Rightarrow \mathtt{Pr}(obs_{i,j}(t) \wedge \phi) \bowtie c \cdot \mathtt{Pr}(obs_{i,j}(t))) \; .$$

A similar transformation applies for more general agent probability comparisons, but we note that linear comparisons may transform to polynomial comparisons: similarly to the discussion of conditional probability in Section 2.5. □

Proposition 5 *With respect to $\mathcal{I}^{\mathtt{spr}}(M)$ for a finite PO-DTMC M, the probability formulas $\mathtt{Pr}_{i,t}(\phi) \bowtie c$ can be reduced to linear comparisons using only terms $\mathtt{Pr}(\psi)$, provided second-order quantifying-in is permitted. Knowledge terms $K_{i,t}$ can be reduced to the universal modality $\Box$, provided second-order quantifying-in is permitted for this modality.*

Proof: Define $\kappa_i(X_1, \ldots, X_{k_i}, t)$ to be the formula

$$\forall t' \leq t(\bigwedge_{j=1\ldots k_i} X_i(t') \Leftrightarrow obs_{i,j}(t'))$$

Intuitively, this says that, up to time t, the second order variables $X_1, \ldots, X_k$ encode the pattern of occurrence of observations of agent i up to time t. The formula

$$\forall X_1, \ldots X_{k_i}(\kappa_i(X_1, \ldots, X_{k_i}, t) \Rightarrow \Box(\kappa_i(X_1, \ldots, X_{k_i}, t) \Rightarrow \phi)$$

is easily seen to be equivalent to $K_{i,t}(\phi)$ in $\mathcal{I}^{\mathtt{clk}}(M)$. Similarly, $\mathtt{Pr}_{i,t}(\phi) \bowtie c$ can be expressed as

$$\begin{aligned}\forall X_1, \ldots X_{k_i}(\quad & \kappa_i(X_1, \ldots, X_{k_i}, t) \Rightarrow \\ & \mathtt{Pr}(\kappa_i(X_1, \ldots, X_{k_i}, t) \wedge \phi) \bowtie c \cdot \mathtt{Pr}(\kappa_i(X_1, \ldots, X_{k_i}, t))) \,.\end{aligned}$$

□

One might wonder whether the knowledge operators can be eliminated entirely using probability, treating $K_i\phi$ as $\mathtt{Pr}_i(\phi) = 1$. This is indeed the case for formulas ϕ in **CTLPK**. The essential reason is that because formulas of **CTLPK** depend at a point (r, m) only on the run prefix $r[0 \ldots m]$, so the possibility that $\neg\phi$ holds on a non-empty set of measure zero does not occur.

Proposition 6 *For all* **CTLPK** *formulas ϕ and PO-DTMC's M and $x \in \{\mathtt{clk}, \mathtt{spr}\}$ we have $\mathcal{I}^x(M) \models K_i\phi \Leftrightarrow \mathtt{Pr}_i(\phi) = 1$.*

Proof: We write T^x_ϕ for $\{(r, m) \in \mathcal{R}^x(M) \times \mathbb{N} \mid \mathcal{I}^x(M), (r, m) \models \phi\}$ If $\mathcal{I}^x(M), (r, m) \models K_i\phi$ then $\mathcal{K}^x_i(r, m) \subseteq T^x_\phi$. Hence by definition, $[\mathtt{Pr}_i(\phi)]_{\mathcal{I}^x(M),(r,m)} = \mu^x_{r,m,i}(\mathcal{K}^x_i(r, m)) = 1$, so $\mathcal{I}^x(M), (r, m) \models \mathtt{Pr}_i(\phi) = 1$.

Now assume that $\mathcal{I}^x(M), (r, m) \models \mathtt{Pr}_i(\phi) = 1$. This implies that $\mu^x_{r,m,i}(\mathcal{K}^x_i(r, m) - T^x_\phi) = 0$. We have to show that then $Z_i = \mathcal{K}_i(r, m) - T^X_\phi(r, m)$ is empty. Assume otherwise. If $(r', m) \in Z_i$ then $\mathcal{I}^x(X), (r', m) \models \neg\phi$. Now in **CTLPK** we can show by induction that if $\mathcal{I}^x(M), (r, m) \models \psi$ for any formula ψ then $\mathcal{I}^x(M), (r', m) \models \psi$ for any run r' satisfying $r'[0 \ldots m] = r[0 \ldots m]$. (We do not actually need to consider all subformulas of ϕ, but can take the base case to be the subformulas of ϕ of the forms $A\psi$, $E\psi$, $K_j\psi$, p and $f(P_1, \ldots, P_n)$ that are not themselves a subformula of a larger formula of one of these types. The claim is straightforward from the semantics for each of these types of formulas, and an easy induction over the boolean operators then yields the result for ψ.) Hence the set $\mathcal{R}_{r[0\ldots m]}$ of runs having the initial prefix $r[0 \ldots m]$ is a subset of Z_i. Thus, $\mu^x_{r,m,i}(Z_i) \geq \mu^x_{r,m,i}(\mathcal{R}_{r[0\ldots m]}) > 0$. This contradicts the assumption that Z_i is a set of measure 0. □

However, this is not the case for formulas $K_i\phi$ where ϕ is an LTL formula. Consider the following Markov Chain. Here we have, at the initial state s, that $\neg K_i(F\neg q)$, because the agent considers it possible that always q (this holds for all choices of observation functions). However, we have $\mathtt{Pr}_i(F\neg q) = 1$, since the only run where $\neg q$ does

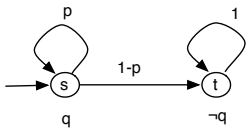

not eventually hold is the run that always remains at s. This run has probability zero. Hence we do not have that $\mathtt{Pr}_i(F\neg q) = 1 \Rightarrow K_i(F\neg q)$ is valid. (The converse form $K_i(\phi) \Rightarrow \mathtt{Pr}_i(\phi) = 1$ does hold for all formulas ϕ of **CTL*****KP**, by the same argument as in the proof of Proposition 6.)

4 Undecidability Results

We can now state the main results of the paper concerning the problem of model checking formulas of (fragments of) the logics **CTL*****KP** and **WMLOKP** in a PO-DTMC M, with respect to an epistemic semantics $x \in \{\mathtt{spr}, \mathtt{clk}\}$. Using the results of Section 3, we also obtain conclusions about extensions of **PMLO** that do not refer to agent probability and knowledge.

For a formula ϕ of **CTL*****KP**, we write $M \models^x \phi$, if $\mathcal{I}^x(M), (r, 0) \models \phi$ for all runs $r \in \mathcal{R}^x(M)$. In the case of **WMLOKP**, we consider sentences, i.e., formulas without free variables, and write $M \models^x \phi$, if $\mathcal{I}^x(M), \tau, r \models \phi$ for all runs $r \in \mathcal{R}^x(M)$ and the empty assignment τ. The model checking problem is to determine, given a PO-DTMC M, a formula ϕ, and semantics $x \in \{\mathtt{clk}, \mathtt{spr}\}$, whether $M \models^x \phi$.

4.1 Background

For comparison with results below, it is worth stating a result from [BRS06] concerning decidability of the fragment **PMLO** of **WMLOKP** that omits knowledge operators $K_{i,t}$ and agent probability terms $\mathtt{Pr}_{i,t}(\phi)$, restricts probability comparisons to the form $\mathtt{Pr}(\phi) \bowtie c$, and prohibits second order quantification to cross into probability terms. Since the structure of agent's local states is irrelevant in this case, we write simply $\mathcal{I}(M)$ for the probabilistic interpreted system corresponding to a PO-DTMC M. To state the result, we define the *parameterized* variant of a formula ϕ of **PMLO** to be the formula $\phi_{x_1,\ldots,x_k}$, in which each subformula of the form $\mathtt{Pr}(\psi) \bowtie c$ is replaced by a formula $\mathtt{Pr}(\psi) \bowtie x_i$, with x_i a fresh variable. We call the resulting formulas the *parameterized formulas of* **PMLO**. For some $\alpha \in \mathbb{Q}^k$, we can then recover the original formula ϕ as the instance ϕ_α obtained from the parameterized variant $\phi_{x_1,\ldots,x_k}$ of ϕ by substituting α_i for x_i for each $i = 1 \ldots k$.

Theorem 7 ([BRS06]) *For each parameterized sentence $\phi_{x_1,\ldots,x_k}$ of* **PMLO***, one can compute for all $\epsilon > 0$ a representation of a set $H_\phi \subset \mathbb{R}^k$ of measure at most ϵ, such that the problem of determining if $\mathcal{I}(M) \models \phi_\alpha$ is decidable for $\alpha \in \mathbb{Q} \setminus H_\phi$.*

Intuitively, the complement of H_ϕ contains the points that are bounded away from limit points of the Markov chain, and comparisons can be decided using convergence properties.

The reason for excluding the set H_ϕ is that the limit point cases seem to require a resolution of problems related to the *Skolem problem* concerning zeros of linear recurrences [Sko34]. A sequence of real numbers $\{u_n\}$ is called a linear recurrence sequence (LRS) of order k if there exist $a_1, \ldots a_k$ with $a_k \neq 0$ such that for all $m \geq 1$,

$$u_{k+m} = a_1 u_{k+m-1} + a_2 u_{k+m-2} + \cdots + a_k u_m \ .$$

We consider the following decision problems associated with a LRS $\{u_n\}$.

1. **Skolem problem.** Does there exist n such that $u_n = 0$?
2. **Positivity problem.** Is it the case that for all n, $u_n \geq 0$?
3. **Ultimate positivity problem.** Does a positive integer N exist such that for all $n \geq N$, $u_n \geq 0$?

We will deal with sequences with rational entries. By clearing denominators the rational version of the above problems can be shown to be polynomially equivalent to similar problems stated using sequences with integer entries. There has been a significant amount of work on these problems [ESPW03], but they have stood unresolved since the 1930's. To date, only low order versions of these problems have been shown to be decidable [HHHK05, OW14, TMS84].

The above problems have an equivalent matrix formulation. A proof of the following can be found in [HHHK05].

Lemma 8 *For a sequence $u_0, u_1, \ldots$, the following are equivalent.*

1. *$\{u_n\}$ is a rational LRS.*
2. *For $n \geq 1$, $u_n = (A^n)_{1k}$ for a square matrix A with rational entries.*
3. *For $n \geq 1$, $u_n = \mathbf{v}^T A^n \mathbf{w}$ where A is a square matrix, and $\mathbf{v}$ and $\mathbf{w}$ are vectors with entries from $\{0, 1\}$.*

In the usual formulation of the Skolem, positivity and ultimate positivity problems, the associated matrices A may contain negative numbers, and numbers not in $[0, 1]$, so are not stochastic matrices. However, [AAOW15] show that these problems can be reduced to a decision problem stated with respect to stochastic matrices:

Lemma 9 *Given an integer $k \times k$ matrix A, one can compute a $k' \times k'$ stochastic matrix B, a length k' stochastic vector $\mathbf{v}$, a length k' vector $\mathbf{w} = (0, \ldots, 0, 1)$ and a constant c such that $(A^n)_{1,k} = 0$ ($(A^n)_{1,k} > 0$) iff $\mathbf{v}^T B^n \mathbf{w} = c$ (respectively, $\mathbf{v}^T B^n \mathbf{w} > c$).*

As noted in [AAOW15], it follows that the logic **PMLO** is able to express the Skolem and positivity problems, using model checking questions of the form

$$M \models \exists t(\mathtt{Pr}(p(t)) = c)$$

and

$$M \models \exists t(\mathtt{Pr}(p(t)) > c)$$

for c a nonzero constant and M a DTMC. (The ultimate positivity problem can also be expressed.) It is worth noting that in the special case of the constant $c = 0$, these model checking questions *are* decidable, as shown in [BRS06]. (Essentially, in this case the problems reduce to graph reachability problems, and the specific probabilities in M are irrelevant.) The transformation from arbitrary matrices A to stochastic matrices B in Lemma 9 requires that the constant 0 of the Skolem problem be replaced by a non-zero constant c.

The above model checking problems of the quantified logic **PMLO** can be seen to be already expressible in the propositional logic **CTL*****KP**, as the problems

$$M' \models^{\mathtt{clk}} \mathbf{AF}(\mathrm{pr}_i(p) = c)$$

$$M' \models^{\mathtt{clk}} \mathbf{AF}(\mathrm{pr}_i(p) > c)$$

$$M' \models^{\mathtt{clk}} \mathbf{AFAG}(\mathrm{pr}_i(p) > c)$$

where we obtain the PO-DTMC M' from the DTMC M by defining $O_i(s) = \perp$ for all states s. That is, agent i is blind, so considers all states reachable at time n to be possible at time n. (We remark that this implies that all the operators A can be exchanged with E without change of meaning of the formulas.) It follows, that with respect to clock semantics, a resolution of the decidability of model checking even these simple formulas of **CTL*****KP** for *all* $c \in [0, 1]$ would imply a resolution of the Skolem problem. In view of the effort already invested in the Skolem problem, this is likely to be highly nontrivial.

4.2 Perfect Recall Semantics

Model checking with respect to the perfect recall semantics is undecidable, even with respect to a very simple fixed formula of the logic **CTL*****KP**, as shown by the following result.

Theorem 10 *The problem of determining, given a PO-DTMC M, if $M \models^{\mathtt{spr}} EF(\mathtt{Pr}_i(p) > c)$, for p an atomic proposition, is undecidable.*

Proof: By reduction from the emptiness problem for probabilistic finite automata [Paz71]. Intuitively, the proof sets up an association between words in the matrix semigroup and sequences of observations of the agent.

A probabilistic finite automaton is a tuple $\mathcal{A} = (Q, \Sigma, \mathbf{v}_0, A, F, \lambda)$, where Q is a finite set of states, Σ is a finite alphabet, $\mathbf{v}_0 : Q \to [0, 1]$ is a probability distribution over states, representing the initial distribution, $A : \Sigma \to (Q \times Q \to [0, 1])$ associates a transition probability matrix $A(a)$ with each letter $a \in \Sigma$, component $F \subseteq Q$ is a set of *final* states, and $\lambda \in (0, 1)$ is a rational number. Each matrix $A(a)$ satisfies $\sum_{t \in S} A(a)(s, t) = 1$ for all $s \in Q$. Let v_F be the column vector indexed by Q with $v_F(s) = 0$ if $s \notin F$ and $v_F(s) = 1$ if $s \in F$. Treating $\mathbf{v}_0$ as a row vector, for each word $w = a_1 \dots a_n \in \Sigma^+$, define $f(w) = \mathbf{v}_0 A(a_1) \dots A(a_n) v_F$. The language accepted by the automaton is defined to be $\mathcal{L}(\mathcal{A}) = \{w \in \Sigma^+ \mid f(w) > \lambda\}$. The emptiness problem for probabilistic finite automata is then, given a probabilistic finite automaton $\mathcal{A}$, to

determine if the language $\mathcal{L}(\mathcal{A})$ is empty. This problem is known to be undecidable [Paz71, CL89].

Given a probabilistic finite automaton $\mathcal{A}$, we construct an interpreted finite PO-DTCM $M_{\mathcal{A}}$ for a single agent (called i rather than 1 to avoid confusion with other numbers) such that $M_{\mathcal{A}} \models^{\mathrm{spr}} EF(\mathtt{Pr}_i(p) > \lambda)$ iff $\mathcal{A}$ is nonempty. This system is defined as follows. We let $N = |\Sigma|$,

1. $S = Q \times \Sigma$,
2. $PI(q, a) = \mu_0(q)/N$,
3. $PT((q, a), (q', b)) = A(b)(q, q')/N$
4. $O_i((q, a)) = a$
5. $p \in \pi((q, a))$ iff $q \in F$.

Note that $\sum_{(q,a)\in S} PI((q, a)) = \sum_{a\in\Sigma} \sum_{q\in Q} \mu_0(q)/N = \sum_{a\in\Sigma} 1/N = 1$, so PI is in fact a distribution. Similarly, for each $(q, a) \in S$, we have $\sum_{(q',b)\in S} PT((q, a), (q', b)) = \sum_{b\in\Sigma} \sum_{q'\in Q} A(b)(q, q')/N = \sum_{b\in\Sigma} 1/N = 1$, so PT is in fact a stochastic matrix.

Note that for each $w = a_1 \dots a_n \in \Sigma^*$ and $a \in \Sigma$, we get a row vector $\mu_{aw} = \mu_0 A(a_1) \dots A(a_n)$ with $\sum_{q\in Q} \mu_{aw}(q) = 1$, which can be understood as a distribution on Q. For each run $r \in \mathcal{R}^{\mathrm{spr}}(M_{\mathcal{A}})$ and $m \geq 0$, we have that that agent i's local state $r_i[0 \dots m]$ at (r, m) is a word in Σ^+. Let $\mathcal{B}(q, m)$ be the set of runs $r \in \mathcal{R}^{\mathrm{spr}}(M_{\mathcal{A}})$ in which $r_e(m) = (q, a)$ for some $a \in \Sigma$. We claim the following about the probability measures $\mu_{r,m,i}$ in the probabilistic interpreted system $\mathcal{I}^{\mathrm{spr}}(M_{\mathcal{A}})$, for each point (r, m) and $q \in Q$:

$$\mu_{r,m,i}(\mathcal{K}_i(r, m)(\mathcal{B}(q, m))) = (\mathbf{v}_0 A(r_i(1)) \dots A(r_i(m)))(q) \; .$$

It is immediate from this that $\mathcal{I}^s pr(M_{\mathcal{A}}), (r, m) \models \mathtt{Pr}_i(p) = c$ where $c = f(r[1 \dots m])$, and the desired result follows.

For the proof of the claim, let μ be the measure on $\mathcal{R}^{\mathrm{spr}}(M_{\mathcal{A}})$ in $\mathbf{Pr}(M_{\mathcal{A}})$. Note first that $\mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r, m))) = 1/N^{m+1}$ for all runs r. The proof of this is by induction on m. Let $r_i(m) = a_0 \dots a_m$. For $m = 0$, we have $\mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r, m))) = \sum_{q\in Q} \mathbf{v}_0(q)/N = 1/N$. For $m = k + 1$,

$$\begin{aligned}
&\mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r, m))) \\
&= \mu(\{r' |\; r'_i[0 \dots k] = a_0 \dots a_k,\; r_e(k) = (q, a_k),\; r_e(k + 1) = (q', a_{k+1})\}) \\
&= \sum_{q\in Q} \sum_{q'\in Q} \mu(\{r' |\; r'_i[0 \dots k] = a_0 \dots a_k,\; r_e(k) = (q, a_k)\}) \cdot PT((q, a_k), (q', a_{k+1}) \\
&= \sum_{q\in Q} \mu(\{r' |\; r'_i[0 \dots k] = a_0 \dots a_k,\; r_e(k) = (q, a_k)\}) \cdot \sum_{q'\in Q} PT((q, a_k), (q', a_{k+1}) \\
&= \sum_{q\in Q} \mu(\{r' |\; r'_i[0 \dots k] = a_0 \dots a_k,\; r_e(k) = (q, a_k)\}) \cdot \sum_{q'\in Q} A(a_{k+1})(q, q')/N \\
&= \sum_{q\in Q} \mu(\{r' |\; r'_i[0 \dots k] = a_0 \dots a_k,\; r_e(k) = (q, a_k)\}) \cdot 1/N \\
&= \mu(\{r' |\; r'_i[0 \dots k] = a_0 \dots a_k\}) \cdot 1/N \\
&= 1/N^{m+1}
\end{aligned}$$

with the last step by induction.

For the main claim, we again proceed by induction on m. For the case $m = 0$, let $r_e(0) = (q, a)$ we have $(r, m) \sim_i (r'm')$ iff $r'_e(0) = (q', a)$ for some q', so $\mathcal{K}_i(r, m) = \{r' \mid r'_e(0) = (q', r_i(0)) \text{ for some } q' \in Q\}$. Hence

$$\begin{aligned}
&\mu_{r,m,i}(\mathcal{K}_i(r,m)(\mathcal{B}(q,m))) \\
&= \mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r,m)(\mathcal{B}(q,m))))/\mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r,m))) \\
&= \mu(\{r' \mid r'_e(0) = (q, r_i(0))\})/\mu(\{r' \mid r'_e(0) = (q', r_i(0)) \text{ for some } q' \in Q\}) \\
&= PI(q, r_i(0))/PI(\{(q', r_i(0)) \mid q' \in Q\}) \\
&= (\mathbf{v}_0(q)/N)/(1/N) \\
&= \mathbf{v}_0(q)
\end{aligned}$$

as required.

For $m = k + 1$, note

$$\begin{aligned}
&\mu_{r,k+1,i}(\mathcal{K}_i(r,k+1)(\mathcal{B}(q,k+1))) \\
&= \mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r,k+1)(\mathcal{B}(q,k+1))))/\mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r,k+1))) \\
&= N^{k+2} \cdot \mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r,k+1)(\mathcal{B}(q,k+1)))
\end{aligned}$$

by the above fact. Now, with $r_i[0 \ldots k+1] = a_0 \ldots a_{k+1}$, we have

$$\begin{aligned}
&\mu(\mathcal{R}^{\mathrm{spr}}(\mathcal{K}_i(r,k+1)(\mathcal{B}(q,k+1)))) \\
&= \mu(\{r' \mid r'_i[0 \ldots k+1] = a_0 \ldots a_k,\ r'_e(k+1) = (q, a_{k+1})\}) \\
&= \sum_{q' \in Q} \mu(\{r' \mid r'_i[0 \ldots k] = a_0 \ldots a_k, r'_e(k) = (q', a_k),\ r'_e(k+1) = (q, a_{k+1})\}) \\
&= \sum_{q' \in Q} \mu(\{r' \mid r'_i[0 \ldots k] = a_0 \ldots a_k, r'_e(k) = (q', a_k)\}) \cdot PT((q', a_k), (q, a_{k+1})) \\
&= \sum_{q' \in Q} \mu(\{r' \mid r'_i[0 \ldots k] = a_0 \ldots a_k, r'_e(k) = (q', a_k)\}) \cdot A(a_{k+1})(q', q)/N \\
&= \sum_{q' \in Q} \mu(\{r' \mid r'_i[0 \ldots k] = a_0 \ldots a_k, r'_e(k) = (q', a_k)\}) \cdot A(a_{k+1})(q', q)/N \\
&= \sum_{q' \in Q} (\mathbf{v}_0 A(a_1) \ldots A(a_m))(q') \cdot A(a_{k+1})(q', q)/N^{k+1} N \qquad \text{(by induction)} \\
&= \mathbf{v}_0 A(a_1) \ldots A(a_m) A(a_{k+1})(q)/N^{k+2}
\end{aligned}$$

Combining this with the previous equation, we have $\mu_{r,k+1,i}(\mathcal{K}_i(r, k+1)(\mathcal{B}(q, k+1))) = \mathbf{v}_0 A(a_1) \ldots A(a_m) A(a_{k+1})(q)$ as required. □

We remark that this result stands in contrast to the situation for model checking the logic of knowledge and time. Write **CLTL*****K** for the logic obtained from **CTL*****KP** by omitting the probability comparison atoms $f(P_1, \ldots, P_k) \bowtie c$. Model checking the

logic **CLTL*****K** with respect to perfect recall, i.e., deciding $M \models^{\mathtt{spr}} \phi$ for M a PO-DTMC and ϕ a formula is decidable [MS99]. (Here, for the semantic structures M, it suffices to replace the initial distribution PI in M by the set $I = \{s \in S \mid PI(s) > 0\}$, and replace the transition distribution function PT in M by the relation R of possibility of transitions between states defined by sRt if $PT(s,t) > 0$. The results in [MS99] use linear time temporal logic as a basis, but, as noted in [MW03], the modality A of the branching time logic CTL^* can be understood as a special case of a knowledge modality: see Proposition 3.)

For probabilistic automata the minimum size of the state space giving undecidability directly stated in the literature appears to be 25 [Hir06]. We remark that the proof of Theorem 10 can also be done by reduction of the following matrix semigroup problem: *given a finite set of matrices of order n, generating a matrix semigroup S, determine whether there is $M \in S$ such that $(M)_{1n} = 0$* [Hal97]. The case of k generators of size $n \times n$ can be reduced to probabilistic automata with $2kn + 1$ states. Recent results on the matrix semigroup problem are given in [CHHN14].

Huang et al [HSZ12] have previously used a reduction from probabilistic automata to show undecidability of an probabilistic epistemic logic with respect to perfect recall. Compared to our simple CTL temporal operators, their logic uses more expressive setting of alternating temporal logic operators.

4.3 Clock Semantics

The undecidability of the perfect recall semantics for such simple formulas suggests that we weaken the epistemic semantics to the clock case. The combination of the translation from **CTL*****KP** to **WMLOKP** (Proposition 3) and Theorem 7 then enables some cases of **CTL*****KP** to be decided. We do not obtain a full decidability result, however, since we face the problem that, with respect to the clock semantics, the formula $AF(\mathtt{Pr}_i(p) = c)$ can express the Skolem problem, so resolving its decidability is a very difficult problem. Rather than attempt to resolve this question, we consider here just how much extra expressiveness is required over the logic of Theorem 7 for us to obtain a definitive *undecidability* result, instead of a decidability result with some excluded and unresolved cases.

Note that one of the restrictions on **PMLO** used in Theorem 7 is that second order quantification should not cross into probability terms. It turns out that this restriction is essential, as shown by the following result.

Theorem 11 *It is is undecidable, given a PO-DTMC M and a formula ϕ of* **WMLOKP** *with linear combinations of probability terms* $\mathtt{Pr}(\phi)$ *and quantifying-in of second-order quantifiers, whether $M \models \phi$.*

Proof: This follows from the fact that, using second order quantifying-in, we can express perfect recall (Proposition 5), and the undecidability of model checking perfect recall (Theorem 10). □

Note that the result refers to $\models$ rather than $\models^{\mathtt{clk}}$, since epistemic operators are not required. This is really a result about a generalization of **PMLO**. One of the

other restrictions in Theorem 7 is that only simple probability comparisons of the form $\mathtt{Pr}(\phi) \bowtie c$ are permitted. More general comparisons of probability terms are needed in applications (see discussion in Section 2.5), so it is of interest to study their impact on decidability. Unfortunately, it turns out to be quite negative. Even the simple case of mixed time polynomial atomic probability formulas is enough for undecidability.

Theorem 12 *There exists a fixed PO-DTMC M with 4 states such that it is undecidable, given a mixed-time polynomial atomic probability formula ψ, whether $M \models \psi$.*

Proof: By reduction from Hilbert's tenth problem, i.e., the problem of determining whether a polynomial with integer coefficients has solutions in the natural numbers. This was shown to be undecidable by Matiyasevich [Mat93].

We show that we can find a stochastic matrix M and a stochastic vector $\mathbf{f}$ such that for each function $f(t) = t \cdot \lambda^t$ and $f(t) = \lambda^t$ with $\lambda = 1/2$, there is a rational vector $\mathbf{g}$ such that $f(t) = \mathbf{f}^T M^t \mathbf{g}$. Given a polynomial $p(n_1, \ldots, n_k)$, we can construct a variant polynomial q' over a larger set of variables, such that an appropriate substitution of such functions $t_i \cdot \lambda^{t_i}$ and λ^{t_i}, for the n_i and the additional variables yields an expression $\lambda^{d_1 t_1 + \ldots + d_k t_k} \cdot p(t_1, \ldots, t_k)$, where the d_i are constants. This has a zero in the $t_1 \ldots t_n$ iff $p(x_1, \ldots, x_n)$ has a zero. It follows that mixed-time polynomial atomic probability formulas can express Hilbert's tenth problem.

Consider the following matrix

$$A = \begin{pmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{4} & \frac{1}{2} & 0 & \frac{1}{4} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{4}$$

It is easily checked that A is stochastic matrix. Expanding by the last row, the characteristic (and minimal) polynomial of A is given by

$$f(x) = (x-1)(x-\frac{1}{2})^2(x-\frac{1}{3})$$

Using Perron's equation [Rom70], which provides a way to express the powers of a matrix using a polynomial, we have

$$\begin{aligned} A^n = {} & 6(A-\frac{1}{2}I)^2(A-\frac{1}{3}I) \; - \; 12(\frac{1}{2})^n(A-I)(A-\frac{1}{3}I) \; - \\ & 12[(\frac{1}{2})^{n-1}n - 4(\frac{1}{2})^n](A-I)(A-\frac{1}{2}I)(A-\frac{1}{3}I) \; - \; 54(\frac{1}{3})^n(A-I)(A-\frac{1}{2}I)^2 \end{aligned} \tag{5}$$

(The reader may verify the correctness of this equation using the fact (see e.g., [Gan59], Ch. V) that if, for a polynomial $p(x)$ and a matrix B, we have, for each eigenvalue λ of B, of multiplicity k, we have $p^{(j)}(\lambda) = 0$ for all $j = 0 \ldots k-1$ (where $p^{(j)}$ is the j-th derivative), then $p(B) = 0$.)

Define the vectors

$$\mathbf{f} = \begin{pmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{pmatrix} \quad \mathbf{g} = \begin{pmatrix} 8/3 \\ 8/3 \\ -8 \\ 0 \end{pmatrix} \quad \mathbf{g}' = 2(A-\frac{1}{2}I)\mathbf{g} = \begin{pmatrix} 0 \\ 4/3 \\ 8/3 \\ 0 \end{pmatrix} \quad \mathbf{g}'' = 2\mathbf{g}' + \mathbf{g} = \begin{pmatrix} 8/3 \\ 0 \\ -8/3 \\ 0 \end{pmatrix}$$

Then $\mathbf{g}'$ is an eigenvector of A with eigenvalue $1/2$, and $(A - \frac{1}{2})^2\mathbf{g} = 0$. Setting $k(x) = (x-1)(x-\frac{1}{3})$, we also have $\mathbf{f}^T k(A)\mathbf{g} = 0$ and $\mathbf{f}^T k(A)\mathbf{g}' = -\frac{1}{12}$. Using equation 5, we obtain

$$\mathbf{f}^T A^n \mathbf{g} = (\frac{1}{2})^n [n-2] \qquad \mathbf{f}^T A^n \mathbf{g}' = (\frac{1}{2})^n \qquad \mathbf{f}^T A^n \mathbf{g}'' = (\frac{1}{2})^n \cdot n \tag{6}$$

for all $n \geq 0$.

We construct a PO-DTMC M for a single agent (called 1), from $\mathbf{f}$ and A, by taking states $S = \{s_1, s_2, s_3, s_4\}$ corresponding to the four dimensions, the initial distribution PI to be given by $\mathbf{f}$, the transition probabilities PT to be given by A, and the observation function defined by $obs_1(s) = \perp$ for all $s \in S$. (That is, the agent is blind.) We define π to interpret the propositions $p_1, \ldots, p_4$ by $p_i \in \pi(s_j)$ iff $i = j$. That is, each p_i holds at state s_i only. Let $\mathcal{I} = \mathcal{I}(M)$.

Corresponding to the vector $\mathbf{g}'$, we define the probability term

$$x(t) = \frac{4}{3}\mathtt{Pr}_{1,t}(p_2) + \frac{8}{3}\mathtt{Pr}_{1,t}(p_3)$$

and corresponding to the vector $\mathbf{g}''$, we define the probability term

$$w(t) = \frac{8}{3}\mathtt{Pr}_{1,t}(p_1) - \frac{8}{3}\mathtt{Pr}_{1,t}(p_3)\ .$$

Then for all runs r of $\mathcal{I}$ and temporal assignments τ such that $\tau(t_i) = n_i$, we have

$$[x(t_i)]_{\mathcal{I},\tau,r} = (\frac{1}{2})^{n_i} \quad \text{and} \quad [w(t_i)]_{\mathcal{I},\tau,r} = (\frac{1}{2})^{n_i} \cdot n_i\ .$$

Now, any diophantine equation

$$p(n_1, n_2, \ldots, n_k) = \sum_{i_1,\ldots,i_k} a_{i_1,\ldots,i_k} n_1^{i_1} n_2^{i_2} \ldots n_k^{i_k} = 0$$

can be expressed as follows. First let d_i be the degree of P in the variable n_i. Corresponding to the monomial $m_{i_1,\ldots,i_k} = n_1^{i_1} n_2^{i_2} \ldots n_k^{i_k}$, define the probability term

$$z_{i_1,\ldots,i_k} = w(t_1)^{i_1} \ldots w(t_k)^{i_k} x(t_1)^{d_1 - i_1} \ldots x(t_k)^{d_k - i_k}\ .$$

Using these terms, we obtain the probability term

$$Z(t_1, \ldots, t_k) = \sum_{i_1,\ldots,i_k} a_{i_1,\ldots,i_k} z_{i_1,\ldots,i_k}\ .$$

Now, for the temporal assignment τ with $\tau(t_i) = n_i$ for $i = 1 \ldots k$, we have

$$[Z(t_1, \ldots, t_k)]_{\mathcal{I},\tau,r} = \left(\frac{1}{2}\right)^{n_1 d_1 + \cdots + n_k d_k} \cdot \sum_{i_1,\ldots,i_k} a_{i_1,\ldots,i_k} n_1^{i_1} n_2^{i_2} \ldots n_k^{i_k}$$

Hence, we have $[Z(t_1, \ldots, t_k)]_{\mathcal{I},\tau,r} = 0$ iff $P(n_1, n_2, \ldots, n_k) = 0$. It follows that $\mathcal{I} \models \exists t_1 \ldots t_k (Z(t_1, \ldots, t_k) = 0)$ iff P has a solution in the natural numbers. □

We remark that the possibility of encoding Hilbert's tenth problem is not immediate from the fact that we are dealing with polynomials, since our polynomials are over *rational* values generated in a very specific way from Markov chains, rather than arbitrary integers. Indeed, there are decidable logics containing polynomials, such as the theory of real closed fields [Tar51].

As noted in Section 2.5, formulas (allowed by Theorem 7) of the form $\mathtt{Pr}(\phi(t_1, \ldots, t_n)) \bowtie c$ can be written as a polynomial of probability expressions, so it is natural to ask whether such formulas also suffice to make the logic undecidable. This does not seem to be the case: the polynomials involved have only positive coefficients. Since Hilbert's tenth problem is trivially decidable for polynomials with only positive coefficients, our proof does not apply to this case.

5 Conclusion

Our results have by no means resolved Skolem's problem, which remains an apparent barrier to resolving the gap between the decidability results of [BRS06] and the undecidability results of the present paper.

However, in work to be presented elsewhere, we show that the results of [BRS06] can be extended both by reducing the set H_ϕ of cases that needs to be excluded to obtain decidability, as well as enhancing the expressiveness to cover epistemic probabilistic terms of the form $\mathtt{Pr}_i(\phi)$, interpreted with respect to the clock semantics.

Acknowledgment: The authors thank Igor Shparlinski and Min Sah for illuminating discussions on Skolem's problem, and Xiaowei Huang for helpful comments on a draft of the paper.

References

[AAOW15] S. Akshay, T. Antonopoulos, J. Ouaknine, and J. Worrell. Reachability problems for Markov chains. *Information Processing Letters*, 115(2):155–158, 2015.

[ACD90] R. Alur, C. Courcoubetis, and D. L. Dill. Model-checking for real-time systems. In *Proc. of the Symposium on Logic in Computer Science*, pages 414–425, 1990.

[BRS06] D. Beauquier, A. M. Rabinovich, and A. Slissenko. A logic of probability with decidable model checking. *J. Log. Comput.*, 16(4):461–487, 2006.

[Buc60] J. R. Buchi. Weak second order arithmetic and finite automata. *Zeitscrift fur mathematische Logic und Grundlagen der Mathematik*, 6:66–92, 1960.

[CE81] E. M. Clarke and E. A. Emerson. The design and synthesis of synchronization skeletons using temporal logic. In *Proc. of the Workshop on Logics of Programs, IBM Watson Research Center, LNCS 131*, 1981.

[CES86] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Trans. Program. Lang. Syst.*, 8(2):244–263, 1986.

[CHHN14] J. Cassaigne, V. Halava, T. Harju, and F. Nicolas. Tighter undecidability bounds for matrix mortality, zero-in-the-corner problems, and more. *arXiv*, abs/1404.0644, 2014.

[CL89] A. Condon and R. J. Lipton. On the complexity of space bounded interactive proofs. In *Proc. of the Symp. on Foundations of Computer Science*, pages 462–467. IEEE, 1989.

[EGM07] K. Engelhardt, P. Gammie, and R. van der Meyden. Model checking knowledge and linear time: PSPACE cases. In *Proc. of the Int. Symp. on Logical Foundations of Computer Science LFCS*, pages 195–211, 2007.

[Eij04] D.J.N. Eijck. Dynamic epistemic modelling. *CWI. Software Engineering [SEN]*, (E 0424):1–112, 2004.

[ESPW03] G. Everest, I. Shparlinski, A. J. van der Poorten, and T. Ward. *Recurrence sequences*. Amer. Math. Soc., 2003.

[FH94] R. Fagin and J. Y. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, 1994.

[FHM90] R. Fagin, J. Y. Halpern, and N. Megiddo. A logic for reasoning about probabilities. *Information and Computation*, 87(1/2):78–128, 1990.

[Gan59] F.R. Gantmacher. *Theory of Matrices*, volume 1. AMS Chelsea Publishing, 1959.

[GM04] P. Gammie and R. van der Meyden. MCK: Model checking the logic of knowledge. In *Proc. Conf. on Computer-Aided Verification, CAV*, pages 479–483, 2004.

[Hal97] V. Halava. Decidable and undecidable problems in matrix theory. Technical Report 127, Turku Centre for Computer Science, University of Turku, Finland, 7 1997.

[Hal03] J. Y. Halpern. *Reasoning about Uncertainty*. MIT Press, Cambridge, MA, USA, 2003.

[HHHK05] V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. Skolem's problem - on the border between decidability and undecidability. Technical Report 683, Turku Centre for Computer Science, University of Turku, Finland, 4 2005.

[Hir06] M. Hirvensalo. Improved undecidability results on the emptiness problem of probabilistic and quantum cut-point languages. Technical Report 769, Turku Centre for Computer Science, University of Turku, Finland, 5 2006.

[HLM11] X. Huang, C. Luo, and R. van der Meyden. Symbolic model checking of probabilistic knowledge. In *Proc. of the Conf. on Theoretical Aspects of Rationality and Knowledge*, pages 177–186, 2011.

[HM10] X. Huang and R. van der Meyden. The complexity of epistemic model checking: Clock semantics and branching time. In *Proc. ECAI 2010 - European Conf. on Artificial Intelligence*, pages 549–554, 2010.

[HSZ12] X. Huang, K. Su, and C. Zhang. Probabilistic alternating-time temporal logic of incomplete information and synchronous perfect recall. In *Proc. AAAI*, 2012.

[HV91] J. Y. Halpern and M. Y. Vardi. Model checking vs. theorem proving: A manifesto. In *Proc. of the Int. Conf. on Principles of Knowledge Representation and Reasoning*, pages 325–334, 1991.

[Kam68] H. W. Kamp. *Tense logic and the theory of linear order*. PhD thesis, University of California, Los Angeles, 1968.

[KNN+08] M. Kacprzak, W. Nabiałek, A. Niewiadomski, W. Penczek, A. Półrola, M. Szreter, B. Woźna, and A. Zbrzezny. Verics 2007 - a model checker for knowledge and real-time. *Fundamenta Informaticae*, 85(1):313–328, 2008.

[KSK76] J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov Chains*. Springer-Verlag, 1976.

[LQR09] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: a model checker for the verification of multi-agent systems. In *Proc. Conf. on Computer-Aided Verification*, pages 682–688, 2009.

[Mat93] Yuri V. Matiyasevich. *Hilbert's Tenth Problem*. MIT Press, 1993.

[MNP08] O. Maler, D. Nickovic, and A. Pnueli. Checking temporal properties of discrete, timed and continuous behaviors. In *Pillars of Computer Science, Essays Dedicated to Boris (Boaz) Trakhtenbrot on the Occasion of His 85th Birthday*, pages 475–505, 2008.

[MS99] R. van der Meyden and N. V. Shilov. Model checking knowledge and time in systems with perfect recall. In *Proc. FST-TCS*, pages 432–445, 1999.

[MW03] R. van der Meyden and K-S. Wong. Complete axiomatizations for reasoning about knowledge and branching time. *Studia Logica*, 75(1):93–123, 2003.

[OW14] J. Ouaknine and J. Worrell. Positivity problems for low-order linear recurrence sequences. In *Proc. ACM-SIAM Symposium on Discrete Algorithms*, pages 366–379, 2014.

[Paz71] A. Paz. *Introduction to probabilistic automata*. Academic Press, 1971.

[RKNP04] J. Rutten, M. Kwiatkowska, G. Norman, and D. Parker. *Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems,* P. Panangaden and F. van Breugel (eds.), volume 23 of *CRM Monograph Series*. American Mathematical Society, 2004.

[Rom70] V. I. Romanovsky. *Discrete Markov Chains*. Wolters-Noordhoff, 1970.

[Sha49] C. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, 1949.

[Sko34] T. Skolem. Ein Verfahren zur Behandlung gewisser exponentialer Gleichungen und diophatischer Gleighungen. In *8de Skand. mat. Kongr. Forth, Stockholm*, 1934.

[Tar51] A. Tarski. *A Decision method for elementary algebra and geometry*. Univ. of California Press, 2nd edition, 1951.

[TMS84] R. Tijdeman, M. Mignotte, and T.N. Shorey. The distance between terms of an algebraic recurrence sequence. *Journal für die reine und angewandte Mathematik*, 349:63–76, 1984.